By Brian Bailey

Noise is a fact of life. Almost everything we do creates noise as a by-product and quite often what is a signal to one party is noise to another.

Noise cannot be eliminated. It must be managed. But is noise becoming a larger issue in chips as the technology nodes get smaller and packaging becomes more complex? For some, the answer is a very strong yes, while for others it does not even appear on the radar.

Noise sources

Step one is to understand the sources of noise. "Noise is becoming more challenging since lower Vdd supply voltages are used in shrinking semiconductor geometries," explains Graham Bell, vice president of marketing for Uniquify, Inc.. "Typical core Vdd changes from 1.2V at 65nm to 0.8V at 14nm, a drop of 33%. Lower Vdd means the ratio of the possible signal to noise (SNR) is smaller."

That isn't the only factor getting worse. "The advent of finFET CMOS nodes has led to increased interconnect density and, combined with the smaller block dimensions, they increase the likelihood of noise coupling either due to electromagnetic or silicon substrate crosstalk," says Magdy Abadir, vice president of corporate marketing at Helic. "Noise coupling generally manifests itself as various forms of jitter. But worst of all, it leads to critical timing analysis issues that are hard to close."

In the past, people asked why they should care about device noise. "In addition to all of the noise effects that we are familiar with, when you go into nanometer and especially finFET, the impact of device noise goes from being a third-order to a first-order effect," says Mick Tegethoff, director of product management, DSM AMS for Mentor, a Siemens Business. "By device noise I specifically mean thermal noise on devices or resistors and flicker noise. Both of those effects become really critical in any kind of analog circuit. Our experience is that anyone designing a Sigma-Delta ADC or any high-performance ADC or PLL synthesizer, if they are not accounting for device noise, they are talking huge risks."

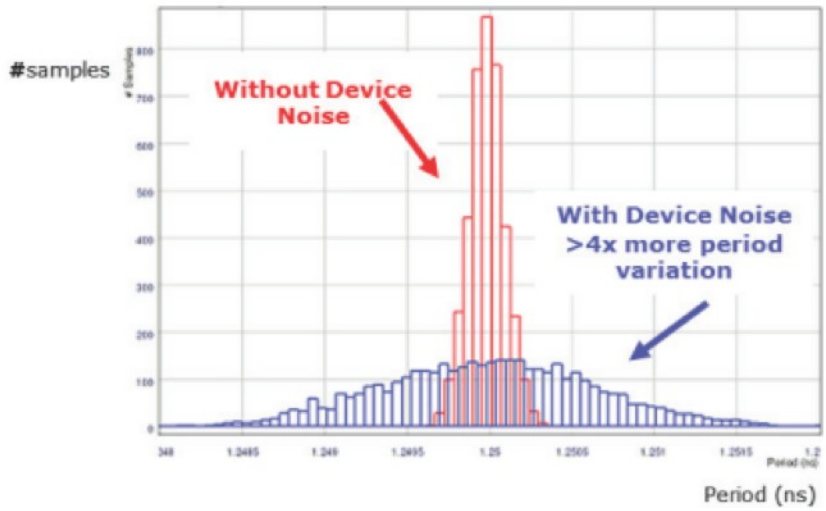

Fig 1: Device noise impact on PLL jitter. Source Mentor.

Coupling is getting more complex. "Noise can propagate through signal

lines, power-delivery network, silicon substrate, interposer substrate layers, package or PCB layers, and the higher frequencies make things just worse," adds Abadir.

"In the old days of 65nm, all of the rules of thumb used to understand noise and to predict noise worked fairly well," says Tegethoff. "When you take device noise and couple that with parasitics from the layout and couple that with the reduced margins of the new technologies, now you have a perfect storm."

Who has to worry?

Not every designer is affected. "Jitter is the biggest noise problem in a PHY," says Navraj Nandra, senior director of marketing for Synopsys' DesignWare Analog & MSIP Solutions Group. "In RF it is phase noise. A PHY is a mixed-signal design and most of our problems are large-signal. The output swing of a transmitter is a few hundred millivolts, and we are interested in the large-signal behavior of that. If you are doing data convertors or RF design, then you are more concerned with the small signal parameters, which are more impacted by noise."

And jitter may not be getting substantially worse. "I am not seeing anything that tells me that jitter is the problem," continues Nandra. "We have been looking at high-speed SerDes from 28nm down to 10nm and 7nm and we are looking at the jitter numbers. Our designers are not saying they have a problem in terms of jitter."

There is more good news. "Foundries have been trying to mitigate the impact of flicker noise by taking advantage of its physical nature, i.e. the charge trapping and de-trapping process," explains Brian Chen, AMS product marketing manager at Synopsys. "They have been perfecting fabrication techniques such as those related to surface preparation and cleaning that result in higher surface quality and thus lower flicker noise."

Other aspects of the smaller geometries are making analysis a lot more complex. "The increased number of dielectrics poses additional challenge as a typical 7nm process node can have as many as 95 different dielectric layers that load any analysis engine with significant overhead," points out Abadir.

Noise analysis

It all starts with the BSIM models. "The BSIM models contain noise modeling," says Tegethoff. "You get this from the foundry and you have to hope that they are doing an accurate model extraction to model thermal noise and flicker noise accurately."

Designs are getting larger. "Tools must offer substantial capacity to operate at the complexity of a low-nm node, large SoC," says Abadir. "The challenge calls for innovative tool architectures using best-in-class theories and solutions that deal with

the full EM field problem. Tools must be able to deal with nm feature sizes and, in particular, the complex dependency of process manufacturing and electrical characteristics on design layout density, known as LDE effects."

The scope of analysis is expanding. "In the past, power noise simulations were done in silos," says Arvind Vel, senior director of applications engineering at ANSYS. "The chip designer, package designer and the board designer would have individual supply noise budgets to meet. As long as they were within their budgets, the overall operation of the chip was guaranteed. With the diminishing noise margins in advanced technology nodes, the silo based margining quickly leads to over-design."

But it is not just about margining. "Shared power and ground domains are important noise pathways that can couple noise and cause issues during operation," continues Vel. "Substrate noise injection is especially important to simulate for RF components integrated along with high speed digital cores."

Even in small devices, noise can be a challenge. "Transient noise it is a random effect so you have to use a lot math to get it right," explains Tegethoff. "We model the random device noise in the transistors every timestep of the simulation and it is a function of the instantaneous bias of the device. You also have to do enough simulation cycles to ensure that you get the full representation. If you are simulating a PLL, you simulate in the time domain but then you have to post process to see it in the frequency domain so that you can see phase noise. That takes a lot of cycles."

Another factor causing increased concern comes from new markets. "With the growth in automotive, functional safety and reliability are becoming a big deal," adds Tegethoff. "And with the variability of the process itself, characterization of analog IP is becoming increasingly important. It puts significant stress on the circuit verification process that continues to require SPICE accurate results and to match silicon but do it a lot faster and with more capacity while doing advanced analysis."

Handling noise

There are multiple places where noise can be tackled. "Most mitigation is done during layout," says Vinod Kariat, corporate vice president of R&D in the Custom IC & PCB Group at Cadence. "But the design itself can be made more robust to noise if noise immunity is considered while doing power optimization."

"Design techniques to mitigate flicker noise include chopping, correlated double sampling, switched biasing etc.," says Synopsys' Chen. "The switched biasing technique, which effectively cycles a MOS transistor between strong inversion and

accumulation, leads to much lower flicker noise in a fashion similar to the recovery phenomenon of the bias-temperature instability (BTI) effect."

The successful application of such techniques requires suitable analysis tools. "HSPICE has recently enabled the calculation of power spectral density directly from transient noise simulation results," continues Chen. "Such capability allows explicit and direct comparison of linear and transient noise and helps validate proper treatment of flicker noise, especially when the designer gets different simulation results from different simulators and needs a golden reference."

For a lot of designers, it will be about balance. "All of the good rules of isolation and layout and managing power continue to be true," says Tegethoff. "Part of the challenge is that the drive for low power tends to work against noise which sometimes requires more current. Good practices continue to apply. If you control your noise margins in the traditional way and control rise times and keep doing the things you did before, it still works for the digital domain. Cell libraries take noise and variation into account such that synthesis does not have to worry about it."

Designs also can adapt. "What is needed for high-speed interfaces that move off-chip is dynamic self-calibration logic (SCL) that can analyze the performance of the channel in real time to maximize throughput despite static and dynamic variations including signal-noise," says Bell. "For communication engineers, this means the eye for inter-signal interference (ISI) is kept wide-open. With dynamic SCL, performance is not left on the table as data-rate is maximized during operation. With the adoption of 2.5D interposer and 3D packaging, there's an even greater need for a real-time approach to deal with static and dynamic thermal variations."

New packaging techniques raise complexity. "As opposed to 2D packaging technologies, where the metal tracks on the SiP substrate are orders of magnitude larger compared to the metal tracks on the silicon chips, 2.5D and 3D IC/SiP integration enables the creation of silicon chip based metal dimensions at either sides (top and back) of the silicon interposer," explains Abadir. "Despite the obvious benefits, this higher integration level facilitates the triggering of several undesired effects, such as electromagnetic crosstalk, due to the higher metal routing density above and below the common silicon interposer and the hundreds or thousands of through-silicon vias implementing die to die or interposer top to back connections. This increases the electromagnetic modeling complexity exponentially."

Conclusion

While noise continues to be a problem for certain classes of design, it is not a problem for everyone. For designs that are impacted by noise, analysis is becoming a

lot more complex. Tools are being stretched to take into account all of the physical effects and the coupling in today's designs is increasing. So far, tool vendors appear to be keeping up with the challenge using a variety of techniques, but designers can also do a lot to try and isolate the problems as much as possible.